BANK OF CANADA
BANQUE DU CANADA

Discussion Paper/Document d'analyse
2011-11

A Model of the EFA Liabilities

by Francisco Rivadeneyra and Oumar Dissou

www.bank-banque-canada.ca

Bank of Canada Discussion Paper 2011-11

December 2011

A Model of the EFA Liabilities

by

Francisco Rivadeneyra and Oumar Dissou

Funds Management and Banking Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
Corresponding author: riva@bankofcanada.ca

ISSN 1914-0568

Acknowledgements

For their comments, we thank Zahir Antia, David Bolder, Pascal Farahmand, Mervin Merkowsky, Miguel Molico, Ron Morrow, Philippe Muller, Tomoaki Nakashima, Cadet Pacouloute, George Pickering, Yuliya Romanyuk, Eric Wolfe and participants of the 2010 Fellowship Exchange Seminar. We also thank Daniel Armali for able research assistance. All errors remain ours.

Abstract

The authors describe the liabilities model of the Exchange Fund Account (EFA). The EFA is managed using an asset-liability matching framework that requires currency and duration matching of both sides of the balance sheet. The model chooses the mix of liabilities across instruments and tenors that maximizes the return of the fund subject to a fixed asset-allocation rule and duration matching. The model considers two types of instruments: cross-currency swaps and global bonds. The main trade-off in the model is the cost advantage of cross-currency swaps relative to global bond issuance. Cross-currency swaps are, on average, a cheaper source of funding, but carry counterparty risk. The model penalizes a skewed maturity profile of liabilities because it carries rollover risks. The model also reports the implied asset-liability gap, which is a function of the total amount of cross-currency swaps.

JEL classification: G32, G12
Bank classification: Foreign reserves management; Debt management

Résumé

Les auteurs décrivent le modèle utilisé pour le choix des passifs du Compte du fonds des changes. La gestion du fonds des changes repose sur un cadre dans lequel les actifs et les passifs au bilan sont appariés étroitement au plan de la devise et de la durée. Le modèle relatif aux passifs aide le gestionnaire à choisir une structure qui, du point de vue tant des passifs retenus que des échéances, maximise le rendement du fonds compte tenu de la règle fixée pour la répartition des actifs et du principe d'appariement de la durée. Deux types d'instruments sont considérés dans le modèle : les swaps de devises et les obligations libellées en monnaies étrangères. Le principal arbitrage concerne l'avantage de coût des swaps sur celles-ci. En effet, les swaps de devises constituent généralement une source de financement meilleur marché que les obligations libellées en monnaies étrangères, mais ils présentent un risque de contrepartie. Par ailleurs, un profil asymétrique des échéances est pénalisé en raison des risques de refinancement qu'il comporte. Le modèle permet en outre de calculer l'écart implicite entre la valeur des actifs et celle des passifs, lequel dépend du montant total des swaps de devises.

Classification JEL : G32, G12
Classification de la Banque : Gestion des réserves de change; Gestion de la dette

1 Introduction

The Canadian official international reserves are held mainly in a government account called the Exchange Fund Account (EFA). The EFA is a portfolio primarily made up of securities denominated in U.S. dollars, euros and yen, special drawing rights (SDRs) and a small holding of gold. The objective of the EFA is to ensure that the government can readily meet its foreign currency obligations (for example, to aid in the control and protection of the value of the Canadian dollar, and to aid in meeting foreign currency payments to the IMF). The strategic objectives of the EFA are to: (i) maintain a high standard of liquidity, (ii) preserve the capital value of the fund and (iii) optimize the net return of the portfolio subject to the achievement of the first two objectives. The EFA is managed using an asset and liability matching (ALM) framework. Under this framework, funds are invested in assets that match, as closely as possible, the characteristics of foreign currency liabilities issued to minimize currency and interest rate risks.

To fund the EFA, reserve managers typically borrow in the form of direct foreign debt (of which global bonds are the most common), the Canada bills program or domestic debt that is subsequently swapped into foreign currency through cross-currency swaps (CCS). To fund the EFA assets means to issue a debt instrument and use the cash proceeds to buy a particular asset that matches the currency and duration of that liability. Generally, CCS debt is less expensive and provides more flexibility in managing the debt maturity profile of the EFA portfolio compared to direct foreign debt instruments. CCS are derivative contracts with private financial institutions by which the Bank of Canada, on behalf of the Government of Canada, exchanges the principal and future interest payments of a liability denominated in Canadian dollars for a liability in one of the three EFA currencies. This form of synthetic debt exposes the government to counterparty risk and introduces volatility to its budgetary position. Counterparty risk arises because of the possibility that the counterparty might not fulfill its obligations when the mark-to-market value of the CCS is in favour of the government. Budgetary risk arises because of the volatility induced by the mark-to-market valuation of CCS. To choose the best funding strategy, reserve managers must account for the cost and risk characteristics of each funding instrument.

This paper describes the liabilities model of the EFA. The model is a quantitative tool to measure the cost and risk trade-off inherent in the funding strategy. Specifically, the strategy is the choice that managers make regarding the funding sources of reserve assets in each of the three currencies in which the EFA holds assets. The key trade-off considered by the model is that CCS are generally cheaper than direct foreign debt but carry counterparty risk. The importance of counterparty risk stems from the possibility of such risks materializing precisely during episodes in which EFA liquidity may be needed.

Other factors are also taken into consideration in the model. The maturity profile that results after a succession of funding decisions is important because it can result in a concentrated maturity profile that carries refinancing risk in periods of market instability. A third consideration is the possible impact of CCS on the government's budgetary position. Given potential changes in accounting rules, movements in the mark-to-market price of derivatives would have to be recognized immediately in the budget of the government. In some situations, the government could find itself recognizing the temporary net liability position of the EFA in the federal budget derived from the mismatch in the value of assets and liabilities. A final factor is the availability of reserve assets in a particular currency and tenor with attractive yield. Given the ALM framework of the

EFA, the choice of tenor of the liabilities determines the tenor (and duration) of the assets, which will affect the total net return of the fund. Therefore, the model restricts the asset-allocation rule in a way that reflects the availability of assets in a particular class, tenor and currency.

For a given currency, the liabilities model solves a constrained optimization problem: to choose the mix of CCS and direct issuance in a variety of tenors that maximizes the expected net return of the EFA subject to: (i) a fixed asset-allocation rule according to the EFA guidelines in that currency; (ii) a penalty on returns imposed when the maturity profile of liabilities exhibits a heavy concentration in the short end; and (iii) an increasing penalty on the funding costs for increasing amounts of CCS funding, to reflect increasing counterparty risk. The model reports, but does not consider in the optimization problem, the implied asset-liability gap, which is the expected difference between the value of assets and liabilities for a given amount of outstanding CCS.

The liabilities model is part of a larger research agenda at the Bank of Canada aimed at solving an integrated asset-and-liabilities model. As opposed to other asset-and-liability matched portfolios, the EFA has no anchor on either side of the balance sheet. For example, pension funds, which traditionally also operate with an ALM framework, take the liabilities as given by the future claims of pensioners. Thus, the problem that pension funds solve is the asset allocation that maximizes returns subject to a given level of risk. In the case of the EFA, both sides of the balance sheet are choice variables; therefore, to solve the problem of liability issuance, the model assumes a fixed asset-allocation rule. This rule anchors the asset side of the balance sheet, thus allowing the model to choose the liabilities that maximize returns given the risks inherent only on the liability side. This rule complies with the guidelines established by the Statement of Investment Policy (SIP).

The rest of the paper is organized as follows. Section 2 describes the EFA in detail, with a particular focus on the liability side of the balance sheet, and explains the cost advantage of CCS over direct issuance. Section 3 describes the model. Section 4 shows the results of the benchmark exercise in U.S. dollars as well as results from a few alternative exercises, including a euro version of the model. We assume that both the U.S. dollar and the euro portfolios are independent of each other, since the size of the EFA and the shares in each currency are decided yearly by EFA managers and, to some extent, are independent of the choices suggested by the two models. Future research will address the asset side of the balance sheet by solving an asset-allocation model with the simultaneous determination of assets, liabilities and currency shares. It is important to stress that, for the purpose of this model, the size of the EFA is an exogenous variable. Section 5 concludes.

2 Description of the EFA and Its Liabilities

2.1 The Exchange Fund Account

The EFA contains the majority of Canada's official international reserves. As of 31 March 2010, the total official international reserves were equivalent to $56.7 billion U.S. dollars, of which the EFA was equivalent to $45.2 billion. The rest is composed of gold, IMF reserves and SDRs.[1] Unlike most other countries, in which the balance sheet of the central bank contains the international reserves, in Canada the reserves are part of

[1]These figures, as well as details on the legal framework of the EFA, can be consulted on the Department of Finance website. The Department of Finance prepares a monthly report of the value and composition of the reserves and the EFA. See http://www.fin.gc.ca/pub/oir-ro-eng.asp. Also see the yearly report on the management of the EFA, Department of Finance Canada (2009).

Table 1: Term Structure of EFA Assets by Currency and Tenor as of 31 March 2010 (US$ billions)

	US$	Euro	Yen	Total
Maturing in 1 year	7.34	2.08	0.09	9.51 (21.0%)
Maturing between 1 and 5 years	7.02	6.43	0.28	13.69 (30.2%)
Maturing after 5 years	10.43	11.58	0.0	22.02 (48.7%)
Total	24.8 (54.9%)	20.0 (44.4%)	0.30 (0.7%)	45.2

the balance sheet of the federal government. The Department of Finance and the Bank of Canada develop and implement the EFA's policies and funding program. As fiscal agent for the government, the Bank executes investment and funding transactions and manages EFA cash flows.

The objective of the EFA, as specified in Part II of the Currency Act, is "to aid in the control and protection of the external value of the monetary unit of Canada." This means that assets in the EFA are managed to provide foreign currency liquidity to the government and to promote orderly conditions for the Canadian dollar in the foreign exchange markets. The strategic objectives of the EFA are to: (i) maintain a high standard of liquidity; (ii) preserve capital value; and (iii) optimize returns. Notice that the high-liquidity and capital-preservation objectives go in the same direction in the risk-return space, given the asset classes in which the EFA is allowed to invest. It is the optimization of returns objective that creates the tension between risk and return in the EFA.

The SIP outlines the types of asset classes, risk exposure limits and other details of the EFA's operations. The permitted asset classes are mainly fixed-income securities issued by governments, their guaranteed entities and supranational entities. The credit ratings of the permitted issuers are constrained to the upper marks of the main rating agencies. At the same time, the SIP details the upper bounds, as a share of the level of reserves, that can be invested by each type of issuer depending on its credit rating. The SIP also establishes a tiered structure of assets to achieve the liquidity goal. The EFA is split into two tiers: a liquidity tier and an investment tier. The liquidity tier itself consists of two subtiers. Liquidity tier I consists of highly rated short-term U.S.-dollar denominated assets, such as U.S. Treasury bills, discount notes and overnight bank deposits, and serves to meet the core liquidity requirements of the EFA. Liquidity tier II is composed of marketable U.S. Treasury, U.S. agency and supranational assets.[2]

Over the past several years, the size of the EFA has steadily increased from US$34.5 billion in 2005 to US$45 billion in 2010. This increase in size is explained by the important role that the EFA plays within the government's Prudential Liquidity Plan. Table 1 describes the composition of the EFA's assets by currency and tenor.

There are several measures of the return of the EFA. The carry is a natural measure given by the ALM framework. The carry measures the interest received on the EFA assets minus the interest paid on the liabilities that fund the assets. If both assets and liabilities are held to maturity, the spread between the yield to maturity of both is equal to the carry. However, this is usually not the case, for reasons that the next section will make clear. Other measures of return, however, could include the effects of interest rate movements on the price of the assets. Also, periods of drastic movements in interest rates can have significant effects on the performance

[2]For more details see http://www.fin.gc.ca/treas/frame/gctrmf06e.pdf.

of the EFA. The model described herein uses the carry as the performance measure to consider the third strategic objective of the EFA: maximizing the return. This is done for simplicity and to avoid, to some extent, making recommendations based on the effects on returns from movements in interest rates during the life of the assets.

The next section describes in detail the liability side of the EFA, and explains the types of funding available to the EFA and their risk and cost characteristics.

2.2 The liabilities

The EFA reserve assets described above are funded by Government of Canada debt, either denominated in Canadian dollars or denominated in one of the three foreign currencies in which the EFA assets are invested (U.S. dollars, euros and yen). Canada uses several debt programs to fund the EFA: Canada bills, global bonds, euro medium-term notes (EMTN), Canada notes and CCS.

The Canada bills is a flexible short-term U.S.-dollar debt program. It is flexible and reliable during periods of turmoil because the placements are of a relatively small size. For example, the Canada bills program provided access to funding during the 2008-09 financial crisis. The term of these debt placements is less than one year.

Moving along the term structure, the program to raise medium-term U.S.-dollar funds is the Canada notes program. Its equivalent in euros is the EMTN program. These programs provide an opportunity to raise funds through the issuance of plain vanilla or structured paper targeted to retail investors through relatively small public issues, or to institutional investors through private placements.

During times of normal market conditions, most of the funding of the EFA assets is done through either global bonds or the CCS program. Global bonds are syndicated placements in foreign currency (U.S. dollars and euros), and are relatively large in size (around 2 to 3 billion). The yield of these bond placements is generally attractive relative to comparable sovereigns, but is subject to market conditions at the time of syndication. Also, given their large size, global bonds have to be planned in advance. To illustrate the funding costs of this program, take the last two global bond placements. During September 2009, the 5-year US$3 billion global bond was issued at a cost equivalent to 3-month LIBOR minus 17 basis points, or a spread of 23 basis points above the 5-year U.S. Treasury benchmark bond. The 10-year placement of 3 billion euros in January 2010 was issued at 20 basis points above the prevailing 10-year German benchmark bund.

CCS are the most flexible and cheapest funding source. CCS are contracts with financial institutions for exchanging the principal and coupon payments of a fixed Canadian-dollar obligation for U.S.-dollar or euro proceeds. CCS are flexible because they are over-the-counter instruments and consequently can be tailored to suit specific requirements of size, maturity dates and coupon flows. CCS transactions conducted by the government are covered by unilateral credit support annexes in favour of the Government of Canada, meaning that the government does not post collateral to offset mark-to-market losses on CCS, but receives collateral on its mark-to-market gains. To mitigate the negative impacts that the unilateral collateral requirements cause to its counterparties, the government includes in its credit clauses a minimum amount of risk that it is willing to tolerate without holding collateral. This threshold is dependent upon the counterparty's credit rating.

To understand why CCS are a cheaper source of funding, it is necessary to explain how CCS contracts are structured. A CCS contract is basically an FX forward contract together with an interest rate swap that transforms a liability in one currency into another while hedging the exchange rate risk. CCS are structured as follows. First a domestic fixed-rate liability is issued. A fixed-to-float interest rate swap converts the fixed coupon payments into floating ones. The relevant floating rate in Canada is the 3-month bankers' acceptance rate. The reason for this is that the market for currency swaps deals in floating rates. The currency swap is the next step and it has two parts: first the forward rate that converts the principal of the domestic liability into U.S. dollars, and second the basis contract that converts the floating Canadian-dollar coupons into U.S.-dollar floating coupons at the prevailing Can$/US$ basis rate. The forward rate is determined by the exchange rate market, and the Can$/US$ basis is determined by the flow of supply and demand for such transactions, which is commonly a function of hedging activities. The 10-year Can$/US$ basis rate has historically been very stable between 10 and 15 basis points, but during the financial crisis of 2008-09 this rate moved drastically, with liquidity and positions varying as the crisis unfolded. Notice that the domestic liability does not need to be placed immediately before the other parts of the CCS contract. In fact, the domestic part of the contract is tapped from liabilities recently placed by the government domestic debt auctions, which follow their independent calendar and policies.

To compare funding costs, we approximate the $\tau-$year Canadian-dollar for the U.S.-dollar CCS cost using:

$$ccs_t^{\tau} = USTzero_t^{\tau} + USDswps_t^{\tau} - CADswps_t^{\tau} - basis_t^{\tau}, \tag{1}$$

where $USTzero$ is the zero-coupon rate in U.S. Treasuries, $USDswps$ is the U.S.-dollar swap spread, $CADswps$ is the Canadian-dollar swap spread, and $basis$ is the Can$/US$ basis rate, all at time t and of tenor τ. This formula does not take into account the forward rate, since it is locked-in, but the net funding cost in dollar terms after exchanging the coupon flows from Canadian into U.S. dollars. The euro CCS involves two additional steps: first a swap from 3-month floating rates to 6-month floating rates, which is the standard for the euro market, and a US$/euro basis swap rate. This is the case because most of the world's foreign exchange derivative transactions are done through the main currency pairs, since they provide more liquidity than, in this case, the direct pair Can$/euro.

The Government of Canada receives attractive quotes of swap rates because it presents very low default risk to its counterparties. This is the main determinant of the low funding costs of CCS to the Government of Canada. Figure 1 shows the estimated historical U.S.-dollar funding costs of 5- and 10-year CCS. This plot shows clearly that the funding costs move with the U.S. Treasury rates. However, the relevant comparison is the *difference* between the funding costs of CCS and global bonds. To estimate this difference, we need a time series of the yield of Canadian U.S.-dollar global bonds. There are few of these bonds outstanding and they trade infrequently in secondary markets; therefore, it is hard to obtain a direct measure of their yield over time. To overcome this limitation, we approximate their yield by a linear combination of U.S. Treasury debt and U.S. agencies. This calculation assumes that Canada can borrow (in U.S. dollars) at a lower rate than agencies like Fannie Mae, and that the volatility of the yield is similar to the U.S. Treasuries. Figure 2 shows the difference between the CCS costs and the global bond costs approximated as $\alpha USTzero + (1-\alpha)GSEzero$, where α is set to 0.8 and $GSEzero$ is the zero-coupon yield of a bond index of U.S. agencies of the appropriate maturity. All data are from Bloomberg.

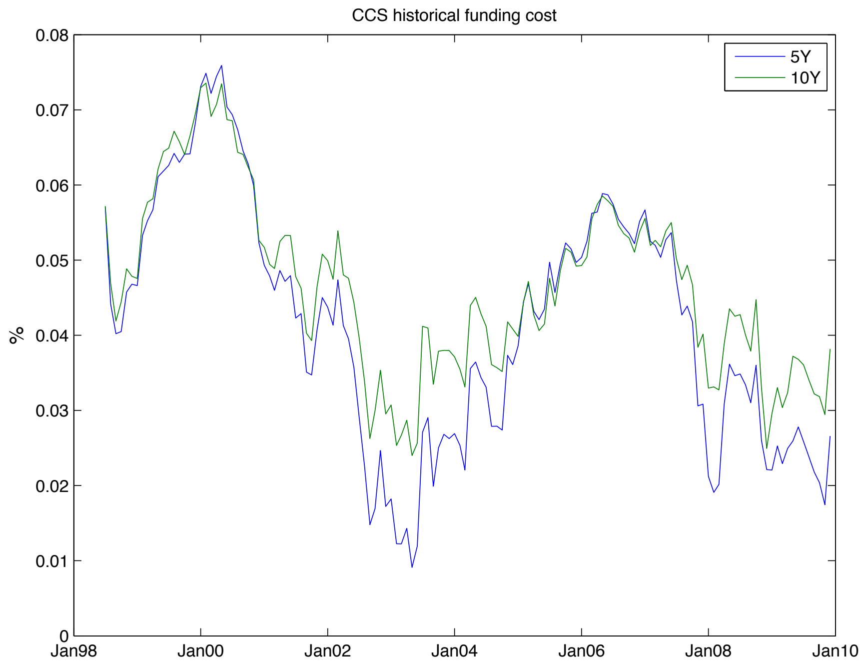

Figure 1: Estimated historical U.S.-dollar funding costs via CCS using equation (1) for the 5- and 10-year tenors.

The cost advantage of CCS comes with credit risk from the counterparties of these contracts. Therefore, the longer the time between entering into a CCS and the fulfillment of the commitment, the greater the credit-risk exposure. The credit exposure stems from the risk of replacing the coupon payments in case of default from the counterparty (Duffie and Singleton 2003, and Duffie and Huang 1996). Although the principal is exchanged at the beginning of the contract, the Government of Canada bears the risk of not finding a replacement of coupon flows at an attractive rate, as well as the risk that the prevailing exchange rate is different from the forward rate that was originally agreed to.

Another source of risk of CCS are changes in interest rates and exchange rates during the life of the contract. This source of risk is relevant for the asset-liability gap, which is the difference in the mark-to-market value of assets and liabilities. In general, the factors that move the mark-to-market valuations of global bonds are tightly linked to the factors that move the valuations of the assets in the EFA balance sheet, since both are fixed-income instruments from highly rated sovereigns. However, factors that move the value of CCS can at times be uncorrelated to the factors that determine the value of EFA assets. An example is the 10-year Can$/US$ basis rate, which, as mentioned above, has been historically stable, but during the 2008-09 financial crisis moved drastically.

There are several variants to the structure of the CCS described above. These differences are important to understand the liabilities model and the mechanics of the EFA's liquidity tier. The CCS described above is called fixed-to-float because it converts a fixed Canadian-dollar liability into a floating U.S.-dollar liability. However, it is also possible to structure a CCS that converts the fixed Canadian-dollar liability into a fixed U.S.-dollar one. The difference between these two lies in swapping the resulting floating U.S.-dollar coupon

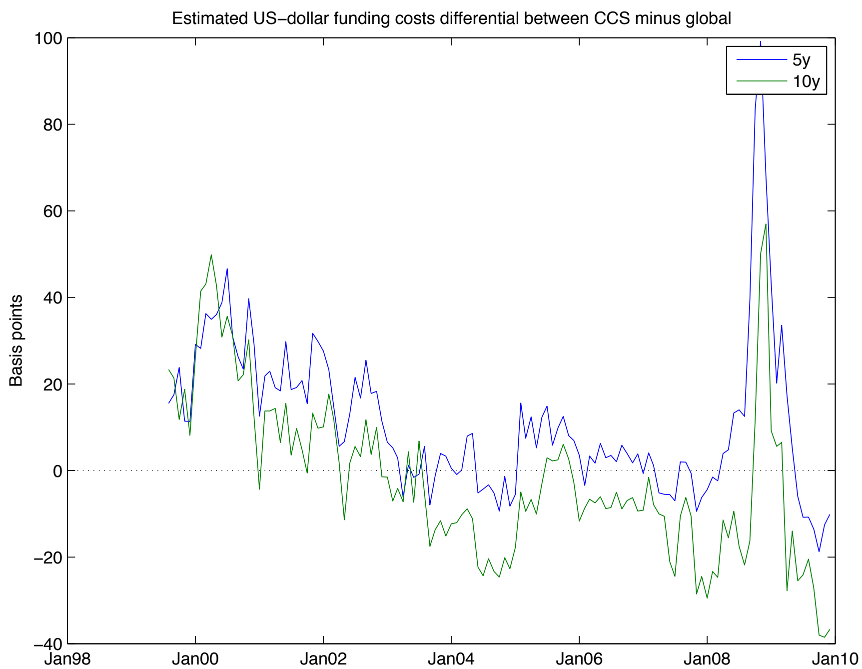

Figure 2: U.S.-dollar funding cost differential between CCS and global bonds. The global funding costs use as a proxy a linear combination of U.S. Treasuries and a U.S. agencies bond index. A number below zero implies that CCS are cheaper than global bonds. During the financial crisis of 2008-09, this difference skyrocketed due to a fall in U.S. yields. Note, however, that these are not transaction data. During the financial crisis, there was very little liquidity in the CCS market.

Table 2: Structure of Funding by Type of Program for Euros and U.S. Dollars as of 31 March 2010. Figures are in billions of U.S. dollars at market exchange rates. Excludes yen liabilities.

	US$	Euro	Total
CCS	19.4	17.2	36.7
Global	3.0	2.8	5.8
Canada bills	2.5	n.a.	2.4
EMTN	n.a.	0.0	0.0
Total	24.8	20.1	44.9

into a fixed U.S.-dollar one via a plain vanilla interest rate swap. Theoretically, the float-to-float and the float-to-fixed are the other two possibilities, but these have not been used regularly in the past several years. Normally, the fixed-to-float and the fixed-to-fixed are used to fund the liquidity tier and the term funding, respectively. Short-term U.S.-dollar assets acquired for the liquidity tier are normally floating-rate instruments, while the investment tier assets are normally fixed-rate assets. This distinction matters for the dynamics of the liability issuance decision because as time passes, some long-term assets become part of the liquidity tier. Tracking these differences over time is important because the model allocates long-term funding only when the liquidity tier funding has been satisfied.

Currently, CCS are the primary funding source of the EFA (Table 2). Only a small portion is currently conducted via direct foreign currency debt (Canada bills and global bonds). Historically, the mix of the EFA funding changed from direct foreign debt to CCS borrowings because of the cost advantage of the latter.

2.3 Literature on asset-liability models for reserves management

The ALM framework is widely used in the management of pension funds and insurance companies. The literature on ALM portfolio management can be split into two camps: academic and practitioners. The academic literature has evolved from work on portfolio allocation. After the work of Markowitz (1952), portfolio theory has been extended from basic mean-variance analysis to more general risk-return analysis in dynamic and stochastic settings (Cochrane 1999 and Wachter 2010), and to include frictions and more realistic settings (Campbell 2006). More recent work on dynamic asset allocation in ALM settings is Detemple and Rindisbacher (2008), which allows the value of the liabilities to vary over time. On the practitioners side, Mulvey and Ziemba (1998) is the main reference. Their approach to solving ALM models is to specify a flexible objective function to accommodate general preferences; the approach deals with uncertainty and its evolution by specifying a reduced number of scenarios instead of sampling a continuum of states of nature, which is computationally challenging. For a broad review of applications, see the handbook of Zenios and Ziemba (2006), which covers ALM models applied to money market funds, insurance and pension funds.

Few other central banks manage their international reserves as Canada does. To some extent, the Bank of England manages its international reserves using an ALM framework. In England, this fund, called the Exchange Equalisation Account, is divided between hedged and unhedged portions. To guide its decisions regarding this account, the Bank of England uses a simple expected-return maximization model that focuses on the likelihood of a call on reserves. Their main aim is to quantify the shadow cost, in terms of forgone returns, of maintaining a high standard of liquidity. Another approach to reserves management is the case study of Colombia by Claessens, Kreuser and Wets (2000). The main contribution of their model, based on the work of Mulvey and Ziemba (1998), is that it specifies the different institutional requirements and objective function of a public institution when the objective of the fund is stabilization and not return maximization.

For a more detailed discussion of these models and issues, see Romanyuk (2010), which surveys the different approaches to modelling ALM portfolios. This paper is mostly concerned with comparing across literatures the measurement of risk and return in ALM portfolios. The dynamic stochastic programming literature is the state of the art in this area, but this approach falls into the category of large-scale models. For example, the famous InnoALM model used by the Austrian subsidiary of Siemens was developed over the course of several years by a relatively large group of researchers (Geyer and Ziemba 2008). The present paper takes a different approach: its aim is to measure in a simple way the main trade-offs in the funding of assets by making a period-by-period issuance decision. The main downside of the non-dynamic approach is that choices today affect the set of possible future choices. Therefore, it is possible, at least theoretically, that cost-saving allocations of liabilities can be achieved if these dynamic effects are taken into consideration.

Conditional on the asset side of the balance sheet, the liabilities model is simply an optimal debt-structure problem. There is a long literature on optimal debt structure in terms of domestic liabilities. Missale (1999) surveys the issues regarding this problem.

3 The Model

3.1 Outline of the model

For a given currency, U.S. dollar or euro, the model chooses a mix of funding across instruments and tenors to maximize the expected net return of the EFA subject to a given level of risk which, on the liabilities side, is given by two measures: the credit risk from counterparties of CCS contracts, and the rollover risk inherent in a skewed maturity profile. The model takes as given the forecasts of funding costs and returns of assets. It also takes as given the size (including future increases) of the EFA portfolio and the maturity profile of liabilities from previous funding decisions.

The funding options are restricted to CCS and direct foreign bonds. For simplicity, the model considers two tenors for each instrument: 5- and 10-year CCS or foreign bonds. The model is solved period-by-period, which implies that financing requirements are met in full every period. Financing requirements in a given period are equal to the liabilities issued in the past that come due plus increments to the EFA, if required. The model has semi-annual frequency, and therefore needs semi-annual forecasts. The horizon of the model is specified by the user, as long as the forecasts of yields and funding costs are supplied.

The period-by-period optimization assumption requires more explanation. It is possible for a dynamic problem to achieve higher net returns if the financing strategies take advantage of movements in the yield of liabilities. However, the Government of Canada does not (traditionally) manage its liabilities actively. Therefore, the funding costs are known almost with certainty at the time of issuing the liabilities. The only uncertainty on the outstanding liabilities is the credit risk of the counterparties of the CCS contract, which are discussed below. Based on this observation, the model assumes that liabilities, once issued, are never bought back, a strategy we call issue-and-redeem. Given this assumption, their yield is riskless (except in the case of a call on reserves or a default of a counterparty of a CCS contract), implying that a dynamic strategy would not be able to reduce the financing costs.

On the asset side, the model assumes a fixed investment rule across three asset classes: sovereigns, agencies and supranationals. This rule is parameterized as portfolio weights for each of the two maturities that the model considers, reflecting the current holdings of the EFA. Furthermore, the model assumes buy-and-hold behaviour for the assets, which implies that the asset manager immediately acquires the asset that matches the duration and desired yield for the given level of risk. In practice, however, this is not the case because the initial purchases after issuing a global bond or a CCS are in assets that are easy to trade. In the case of U.S.-dollar proceeds, these assets are U.S. Treasuries; in the case of euro proceeds, these assets are German bunds. These assets, when purchased initially, have lower yields than the average asset held in the EFA. In practice, the asset positions shift over time toward higher-yielding assets as they become available in the markets.

The measurement of risk on the liability side of the balance sheet is given by two penalty functions: one for credit risks and another for the maturity profile. The penalty for CCS is an increasing function of the issuance amount of CCS. The logic for the increasing shape is that CCS credit risk from default increases in the total outstanding amounts of CCS contracts in the EFA. As the amount of CCS grows, in risk-adjusted terms, the costs of issuance increase, with the increase coming from the implied replacement costs in the case of default. The replacement cost function is parameterized with estimates from the Bank of Canada's Financial Risk Office. Section 3.4 discusses in detail the parameterization of this function. Another interpretation of the increasing cost is that, as outstanding amounts of CCS increase, the cost of replacement increases, because the likelihood of finding counterparties to replace a CCS contract from a defaulting counterparty decreases as defaults can be correlated.

The model considers the liquidity tier to comply with the SIP requirements. The model separates short-term funding requirements from term funding decisions. The total amount of term funding is equal to the total amount of liabilities coming due in a period minus the total amount of short-term funding. Short-term funding for liquidity tier I is allocated to 1-year bonds and fixed-to-floating CCS.[3]

The model also reports a measurement of the expected asset-liability gap. This measure is computed using a simple OLS regression that relates the historical observed asset-liability gap to the components of funding costs in equation (1) and the amount of outstanding CCS. The intuition behind the regression is that the volatility of the mismatch between the value of assets and liabilities has a price component and a quantity component. Most of the variability is explained by the price component that comes from movements in the swap spreads and the Can$/US$ basis, when interacted with the amount of CCS.

With this setting, the model abstracts from the effects of currency diversification. Solving one model for each currency implies that there are no interactions between the returns of the assets in either currency. Further developments will take into account the joint determination of the liabilities in euros and in U.S. dollars.

[3] See section 2.1 for details on the structure of the liquidity tier.

3.2 Details of the model

The problem for each currency can be cast mathematically as a return maximization problem subject to a set of restrictions and bounds. Given the assumptions on the penalties described above, the problem is a convex optimization problem with a unique solution. The bounds are the restrictions that market conditions place on volumes of issuance and the relation that prices have with volumes.

The notation is the following. Let $q_t^j = (q_t^{j,5y}, q_t^{j,10y})'$ be the 1×2 vector of time t issuance in 5- and 10-year tenors of instrument $j = \{b, c\}$, where c is CCS and b is global bonds. Let $a(q_t^c, q_t^b)$ be the asset-allocation rule which is a function that maps the value of the issuance (either in US$ or euros) to the portfolio holdings (at par value) of sovereigns $\left(a_t^{SOV}\right)$, agencies $\left(a_t^{GSE}\right)$ and supranationals $\left(a_t^{SUP}\right)$ for each tenor in the given currency. The rule is the following linear mapping:

$$
\begin{aligned}
a^{\tau}(q_t^c, q_t^b) &= \begin{bmatrix} a_t^{\tau,SOV} \\ a_t^{\tau,GSE} \\ a_t^{\tau,SUP} \end{bmatrix} \\
&= \begin{bmatrix} \alpha^{\tau} \\ \beta^{\tau} \\ \gamma^{\tau} \end{bmatrix} \times (q_t^{c,\tau} + q_t^{b,\tau})
\end{aligned}
$$

where α^{τ}, β^{τ} and γ^{τ} are the τ-year tenor portfolio weights. Note that for each τ we need the restriction $\alpha^{\tau} + \beta^{\tau} + \gamma^{\tau} = 1$. Note also that for each currency (US$ or euros) the portfolio weights may be different, depending on the availability of fixed-income products and the revealed preference of the asset manager.

Again, for each currency, denote the τ-year returns as R_t^{τ}, which is a 3×1 vector for sovereigns, agencies and supranationals. These returns are yields to maturity, and are therefore assumed to be known at time t. Let S_t^j be a 2×1 vector of funding costs for instrument j at time t, with one entry for each of the two maturities. Finally, let $p(q_t^j; j)$ be the penalty function that incorporates replacement and rollover risks. Then the problem for a given currency is to choose quantities q_t^c and q_t^b to solve

$$
\begin{aligned}
\max_{\{q_t^c, q_t^b\}} \quad & \sum_{t=1}^{T} \left[\underbrace{a^{5y}(q_t^c, q_t^b) \times R_t^{5y} + a^{10y}(q_t^c, q_t^b) \times R_t^{10y}}_{\text{return calculation}} - \underbrace{q_t^c \times S_t^c - q_t^b \times S_t^b}_{\text{cost calculation}} - \underbrace{p(q_t^c; c) - p(q_t^b; b)}_{\text{penalties}} \right], \\
\text{s.t.} \quad & \\
\text{budget constraint} \quad & 1 \cdot q_t^c + 1 \cdot q_t^b = x_t + r_t, \\
\text{law of motion} \quad & y_t + x_t = q_{-1}, \\
\text{mkt limits} \quad & \underline{q}^c \le q_t^c \le \bar{q}^c, \\
\text{mkt limits} \quad & \underline{q}^b \le q_t^b \le \bar{q}^b,
\end{aligned}
$$

where x_t is the total amount of term funding, r_t is the amount of increase in the size of the EFA in period t and y_t is the total amount of liquidity tier 1-year funding. The quantity q_{-1} is the total amount of funding coming due from previous funding decisions. It is calculated from the maturity profile and includes the short-term funding from the liquidity tier plus quantities of term funding that may come due in a given period. The last two inequality constraints are market limits on issuance. The lower bar variables are the lower bounds and the upper bar variables are the upper bounds on issuance placed by the size of the market for global bonds and CCS for each period.

Note that the objective function has no time discounting. This could be added if the aim in using the model was to compute the present value of the expected net returns of the EFA. However, the aim is to determine the period-by-period mix of funding. In this case, the time discount factor drops from the first-order condition. The intertemporal linkages of the model are via the "law-of-motion" equation, in which previous funding decisions change the maturity profile of the EFA.

The penalty function is parameterized as the addition of two exponential functions: one for the replacement costs from counterparty risks and a second for the penalty from rollover risks from a skewed maturity profile. The replacement cost penalty is a function of the instrument and its tenor, while the penalty for the skewed maturity profile is only a function of the tenor and applies equally to both instruments. For CCS, the penalty function is

$$p(q_t^c; c) = \left[\left(1 + q_t^{c,5y}\right)^{\theta^{5y}} + \left(1 + q_t^{c,10y}\right)^{\theta^{10y}} + \left(1 + q_t^{c,5y}\right)^{\eta^{5y}} + \left(1 + q_t^{c,10y}\right)^{\eta^{10y}} - 4 \right] / \kappa,$$

where $\theta^{5y}, \theta^{10y}, \eta^{5y}, \eta^{10y} > 0$, and $\kappa = 1,000$. For global bonds the penalty function is given by

$$p(q_t^b; b) = \left[\left(1 + q_t^{b,5y}\right)^{\eta^{5y}} + \left(1 + q_t^{b,10y}\right)^{\eta^{10y}} - 2 \right] / \kappa,$$

meaning that for global bonds there is no replacement cost penalty, since there is no counterparty risk in these instruments.

The natural units of the model are in billions of dollars: therefore, the penalty function adds one unit to each instrument in each tenor to avoid switching sign when $q_t^{j,\tau}$ is below one billion. This scalar is subtracted after the exponential, so the penalty starts from zero for zero issuance. Finally, κ scales the costs into basis points per billion. Importantly, penalties are a function of quantities, because these costs depend on the amounts of outstanding liabilities.

The timing of the model is important. The model assumes that all amounts are due at the beginning of each period. Issuance in the current period (which equal at least all amounts due) are issued instantaneously also at the beginning of the period. Therefore, new issuance will come due at the beginning of the period of its maturity. For example, say \$1 billion comes due in the first semester of 2010 which will be split evenly between liquidity tier and term funding. Then the \$0.5 billion issuance of 1-year bonds will come due in the first semester of 2011 and \$0.5 billion of 10-year bonds will come due in the first semester of 2020.

The treatment of the liquidity tier requires a more detailed explanation. According to the SIP, a minimum of 10 per cent of the EFA assets must be held in short-term U.S.-dollar-denominated assets. These assets are *part* of the liquidity tier, called liquidity tier 1. This only applies to the U.S.-dollar liabilities version of the model. The *requirements* of the liquidity tier, $\bar{y}_t$, are equal to

$$\bar{y}_t = \sum_{i=1}^{\tau} (q_{t,i}^c + q_{t,i}^b) \cdot \phi_t,$$

where $q_{t,i}^j$ is the par value at time t of the bonds and CCS with maturity i, and τ is the longest maturity available in the EFA. ϕ_t is equal to 10 per cent divided by the share of U.S. dollars in the whole EFA. The requirements will change over time since the size of the EFA as well as the share ϕ_t varies over time, depending on exchange rate movements. Note that the requirements and the *actual* liquidity tier are not always equal. The amount of the liquidity tier in period t is given by

$$y_t = \max\left[\sum_{i=1}^{\tau'}(q_{t,i}^c + q_{t,i}^b), \bar{y}_t\right],$$

where τ' is the 1-year maturity. When the amount of total short-term funding is below the requirements, the model allocates automatically the shortfall of liquidity tier in that period. Moreover, the calculation of the actual amounts of short-term funding should also include the fixed-for-floating CCS, which serve the purpose of funding short-term assets even if, in some cases, they are not due within one year. For simplicity, the notation does not take them into consideration, although the model does.

The first-order conditions of the model are useful to understand the mechanics of the model. The first-order condition with respect to the 5-year CCS is

$$\frac{\partial a^{5y}(q_t^c, q_t^b)}{\partial q_t^{c,5y}} \times R_t^{5y} - S_t^{c,5y} - \frac{\partial p(q_t^c; c)}{\partial q_t^{c,5y}} = \lambda_t,$$

where λ_t is the period t Lagrange multiplier of the resource constraint. The condition states that excess returns minus replacement and concentration costs should be equal across each funding source at a given date t. Notice that the difference between the condition of CCS and bonds is that θ terms do not appear in the latter. The condition for CCS adds both penalty parameters.

3.3 Algorithm

The algorithm that solves the maximization problem has two elements: first, the exogenous forecasts of returns and funding costs, and second the optimizer routine. Figure 3 shows graphically the mechanics of the model. The structure of the model and the exogenous elements are fed into the optimizer routine, which numerically solves for the quantities of funding of CCS and global bonds. The solution of the model is unique because the surface being optimized is continuous and strictly concave in all directions and the bounds provide a restricted search area.

The intertemporal linkages of the model are via the maturity profile because the current choices change the future required amounts of issuance in a given period. The algorithm uses the maturity profile every period as the state variable and is updated every period to reflect the period's redemptions and new issuance.

3.4 Calibration of parameters

To solve the model, the optimizer requires the parameters of the penalty functions as well as the bounds on issuance and the asset-allocation rule. The following are the bounds on quantities issued in a given period for each currency (US$ and euro):

1. An upper bound on the volume of global issuance in the 5- and 10-year tenors of 2 billion per currency per semester. The lower bound is zero.

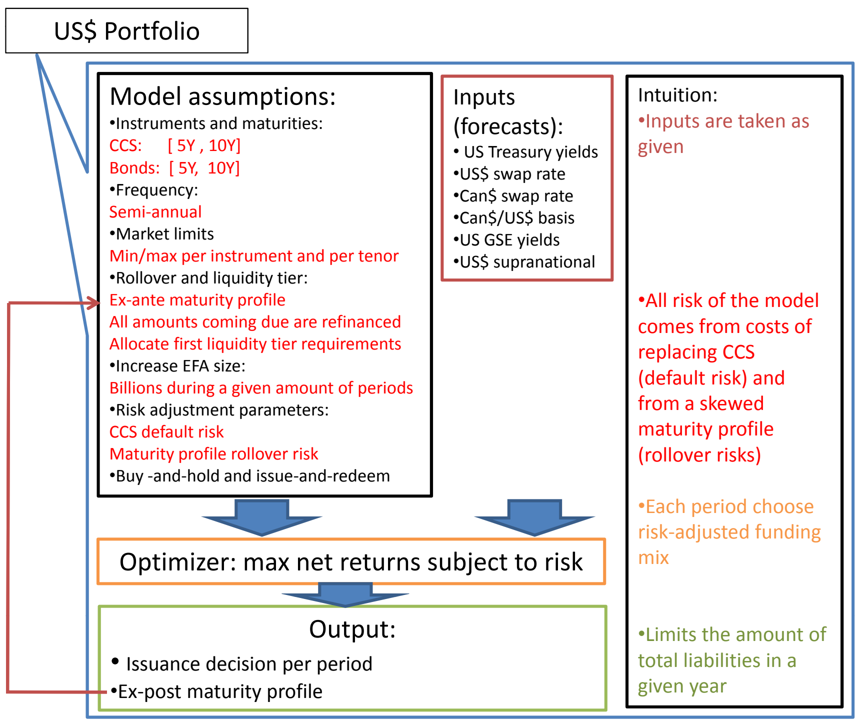

Figure 3: Algorithm for the U.S.-dollar liabilities model. The forecasts of returns and funding costs are exogenous, as well as the parameters of the function of penalties for replacement costs and rollover risks. Also, the bounds given by market limits are exogenous. All these elements are fed into the optimizer, which calculates the funding mix and keeps track of the next period's maturity profile.

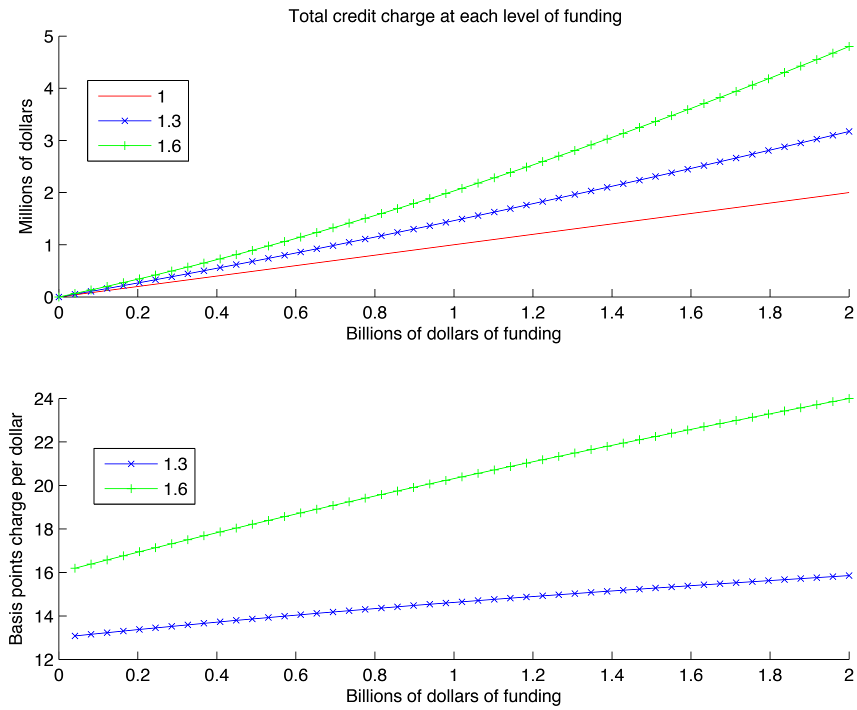

Figure 4: Parameterization of the replacement cost function. The first panel plots the replacement cost as a function of the level of funding for three levels of the parameter θ: 1.6, 1.3 and 1. The second panel plots the ratio of replacement costs divided by the level of funding. This plot shows that, in the margin, a larger level of funding implies a larger replacement cost. The rate of increase is a function of the exponential parameter in the replacement cost function. On average, the replacement costs in the relevant range of 1 to 2 billion are between 15 and 20 basis points. Note that, to increase in the margin, the parameter θ has to be larger than 1.

2. An upper bound on the volume of CCS of 4 billion per currency per semester, which includes 2 billion on each of the maturities. The lower bound is zero.

Another possibility is to consider a non-linear lower bound. That is, given the syndication cost associated with the issuance of global bonds, its is unlikely to issue less than, say, \$500 million. Therefore, the model could consider issuing global bonds if and only if the optimal amount is larger than this bound.

The replacement cost function is calibrated to give penalties between 15 and 20 basis points per dollar (euros) for issuance amounts between 1 to 2 billion dollars (euros). Figure 4 shows the dollar amount of the penalty as well as the per-dollar amounts of charges in basis points for three parameters of θ. The magnitude of these charges were suggested by the Financial Risk Office of the Bank of Canada. For the penalty to be increasing in the margin, θ has to be larger than 1. The maturity profile penalty is also calibrated to compensate for the lower net return of the long-term funding. This is discussed further in the next section. The charge ranges from 10 to 20 basis points in the range from 1 to 2 billion dollars of funding.

The asset-allocation rules will be discussed in section 4.1.

3.5 Asset-liability gap regression

Finally, to report the implied asset-liability gap from the future funding decisions, the model calculates out-of-sample forecasts using the regression described in Table 3. The regression calculates the implied asset-liability gap using the total amount of CCS liabilities outstanding in every period from prior funding decisions and the determinants of funding costs of CCS. Specifically, we regress the changes in the monthly asset-liability gap against the funding cost determinants in equation (1), the level of U.S. Treasury yields, the U.S.-dollar swap spread, the Canadian dollar swap spread and the Can$/US$ basis rate, all interacted with the period's outstanding CCS liabilities. We also use two dummies for outliers in December 2006 and January 2006. The regression result shows that the level of interest rates given by the 10-year U.S. Treasury yield is not important for explaining the changes in the asset-liability gap. This is expected because both sides of the balance sheet co-move with changes in U.S. yields and the U.S.-dollar swap spread. The main determinants for the gap are the changes in the Can$ swap spread and the Can$/US$ basis rate. This is also expected because the factors affecting the valuations from global bonds and assets are very similar, thus leaving most of the gap variation to be explained by the determinants of CCS that do not affect the asset side of the balance sheet.

The out-of-sample forecasts are then calculated using the regression coefficients of the asset-liability gap changes; then level of asset-liability gap is computed using the previous period's gap implied by the funding decisions of the model. The initial level of the gap is the one observed in January 2010.

Table 3: Asset-liability gap regression using monthly frequency data from January 2002 to December 2009. Newey-West standard errors in parentheses. Constant and dummies for outliers are omitted. The variables are the 10-year U.S. Treasury zero-coupon yield, $UST10y$; the 10-year U.S.-dollar swap spread, $USDswsp10y$; the 10-year Canadian-dollar swap spread, $CADswsp10y$; and the Can$/US$ basis rate, $basis$; all interacted by the value of outstanding CCS liabilities.

Independent Var	$D(ALgap)$
$D(UST10y \times CCS)$	0.0178 (0.2697)
$D(USDswsp10y \times CCS)$	0.4393 (1.0113)
$D(CADswsp10y \times CCS)$	-3.1189 (0.8017)
$D(basis \times CCS)$	-4.6786 (0.9035)
$R^2 adj$	0.6852
Number of obs.	95

4 Results and Comparative Static Exercises

This section describes a series of exercises to illustrate the intuition of the model. We use the actual maturity profile data as of 2010 and run two sets of exercises: the first without any penalties and the second using the calibration of penalties discussed above. We also run another set of exercises with a theoretical perfectly smooth maturity profile, to show the implications of each of the elements of the model. These exercises

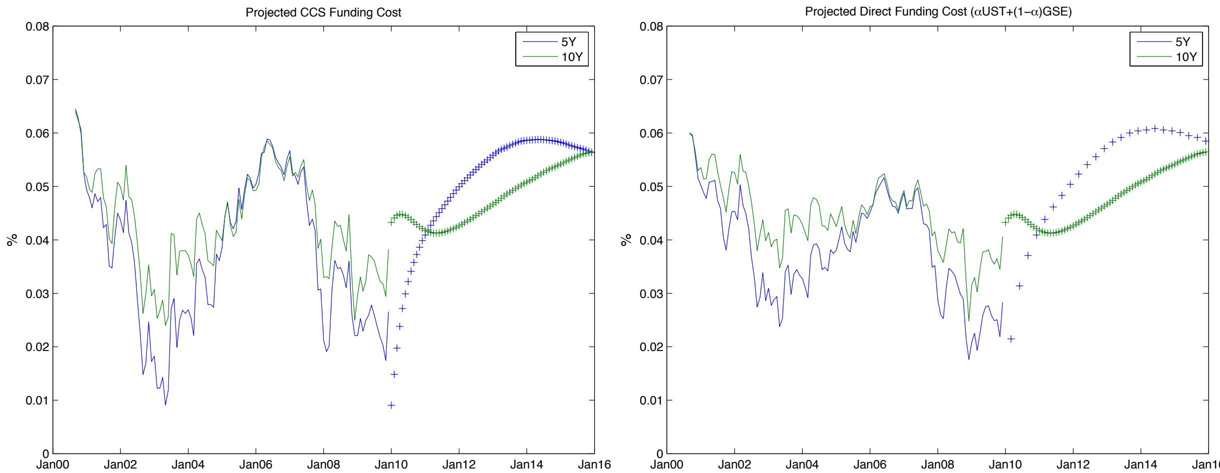

Figure 5: Forecast of CCS and of global bonds funding costs. The left panel shows CCS costs computed using equation (1). The right panel shows the forecast of global bonds funding costs computed as the linear combination of U.S. Treasuries and U.S. agencies using a 0.7 weight for U.S. Treasuries. Agencies themselves are the government yields plus the unconditional mean observed in the past decade. See the text for the values of these spreads by tenor.

are comparative static exercises and are meant to illustrate the mechanics of the model while varying the calibration of the replacement costs, rollover risks and size of the EFA. To perform these exercises, we need asset returns and funding cost forecasts. The following subsections describe the model behind the forecasts of the funding costs and asset returns.

4.1 Asset returns and funding costs

The elements needed to forecast the funding costs of the U.S.-dollar part of the EFA are the U.S. Treasury zero-coupon yields, the U.S.-dollar swap rate, the Can$ swap rate and the Can$/US$ basis rate. CCS funding costs are given by equation (1), in which the U.S.-dollar swap spread is equal to the swap rate minus the U.S. Treasury yield. The global bond costs are simply a linear combination of U.S. Treasuries and the spread of U.S. agency yields.

The U.S. Treasury yield curve is forecast using a modified version of the domestic debt model of Canada (Bolder 2006), which itself is based on a Nelson-Siegel term-structure model. In this version of the model, the main determinants of yields are past yields and the future fiscal position of the United States published by the Congressional Budget Office. The U.S. Treasury zero-coupon curve used to estimate the model is the one provided by the Federal Reserve on its website. The rest of the time series are from Bloomberg. The left panel of Figure 5 plots the CCS historical and forecast funding costs up to December 2015.

We assume that funding costs of global bonds in U.S.-dollars are a linear combination of Treasury funding costs and U.S. agencies. The forecasts of agencies are equal to the U.S. Treasury curve plus the unconditional mean spread between the U.S. Treasury curve and the agencies' yield index from Bloomberg. The spread averages are 0.479 per cent, or almost 48 basis points, for the 5-year tenor, and 0.706 per cent, or 70 basis

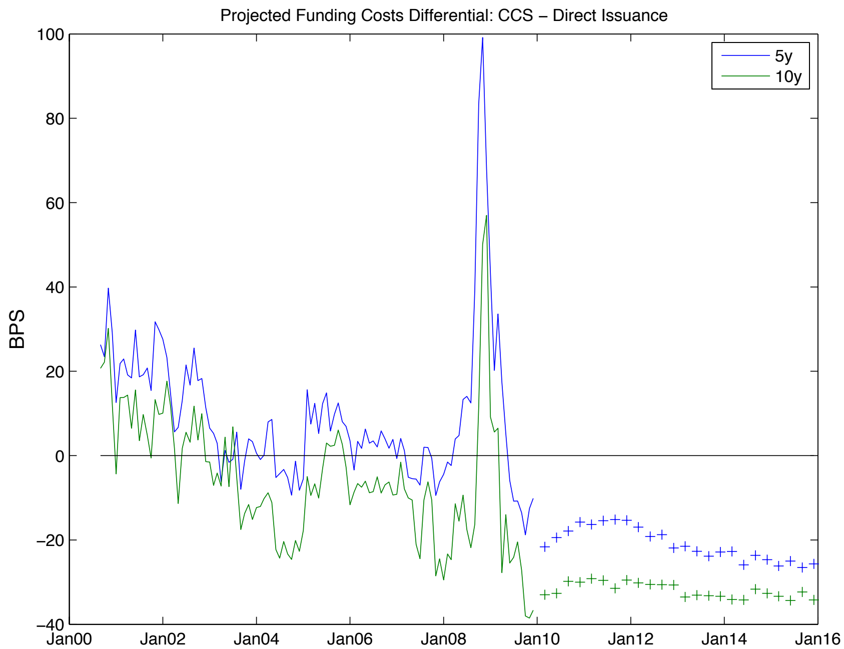

Figure 6: Forecast of cost differential between CCS and global bonds. A number below zero implies that issuing CCS is cheaper on a cost basis than global bonds. By this measure, the 10-year CCS bond is the cheapest funding source. On a net return basis, this is reversed because of the low return of 10-year assets.

points, for the 10-year tenor. Then the funding costs of global bonds in U.S. dollars are obtained by adding a fraction of the agencies' spread to the U.S. Treasury yields. The right panel of Figure 5 plots the estimated global bond historical and forecast funding costs up to December 2015.

Figure 6 shows the forecast of the difference in funding costs between CCS and global bonds for both maturities considered in the model. Notice that the forecasts predict that 10-year CCS are the most advantageous form of funding. This, however, hides the fact that returns of 10-year assets are lower than that of 5-year assets, because of the composition of both asset portfolios.

The asset classes considered in the U.S.-dollar model are sovereigns issuing in U.S. dollars (mostly U.S. Treasuries), agency debt in U.S. dollars (also mostly U.S. agencies) and supranationals (for example, debt from EIB, World Bank and others). To compute the supranationals, we add the unconditional mean spread between the U.S. yields and the Bloomberg index of supranational debt in U.S. dollars. The unconditional mean spread for the 5-year tenor is 0.341 per cent, or 34 basis points, and for the 10-year tenor it is equal to 0.417 per cent or 41 basis points. According to the data, supranational debt in U.S. dollars has a lower yield than U.S. agency debt. The euro version of the model is computed in a similar fashion, using instead the unconditional mean spread between the yield of German bunds and the Bloomberg index of supranational debt in euros.

Combining the asset returns, the parameters $(\alpha^\tau, \beta^\tau, \gamma^\tau)$ for the asset allocation, we can compute per-dollar net returns without considering the penalties for counterparty or rollover risks. This can be done because the net return does not depend on quantities outstanding or issued. On the other hand, the penalties depend on the quantities.

The measure of net returns helps clarify the source of returns, because although 10-year CCS are cheap to issue, there are a very limited amount of 10-year assets with an attractive yield within the SIP guidelines. This is the case because the size of supranational placements, which offer an attractive risk-adjusted yield, are relatively small and are hard to come by.

The asset-allocation rules are given by historical asset holdings in the EFA for each of the two tenors. These parameters are rough estimates and can be modified; however, they are good first-order approximations. For the U.S.-dollar 5-year tenor, the portfolio weights are

$$\begin{pmatrix} \alpha^{5y} & \beta^{5y} & \gamma^{5y} \end{pmatrix} = \begin{pmatrix} 0.5 & 0.25 & 0.25 \end{pmatrix},$$

and for the U.S.-dollar 10-year tenor the weights are

$$\begin{pmatrix} \alpha^{10y} & \beta^{10y} & \gamma^{10y} \end{pmatrix} = \begin{pmatrix} 1.0 & 0.0 & 0.0 \end{pmatrix}.$$

These weights reflect the argument that in the 10-year tenor, it is difficult to get supranationals and agencies in U.S. dollars at an attractive yield. Figure 7 shows the time series of the computed net return over the forecast horizon using the funding costs and the asset returns combined with the asset-allocation rules. As mentioned before, the 10-year tenor does reduce its cost-effectiveness because most assets tend to be placed in low-yielding U.S. Treasuries. From a net return perspective, without considering any penalties, the best funding source is the 5-year CCS.

4.2 U.S.-dollar unrestricted and calibrated exercises

The next three exercises use the actual U.S.-dollar maturity profile and our forecasts of asset returns and funding costs. The difference between the three exercises is the penalties. The first exercise is without any penalty for credit risk or maturity profile. The second exercise adds the credit charge and the third exercise adds the maturity profile. To understand the pattern of issuance over time it is instructive to see the actual maturity profile of U.S.-dollar liabilities. Figure 8 shows the maturity profile of U.S.-dollar liabilities at the end of 2009 split into three categories: fixed-to-float CCS, fixed-to-fixed CCS and global plus Canada bills. These exercises use the liquidity tier rule to allocate the appropriate short-term funding; the remaining amounts are the ones shown in the following plots.

The result of the first exercise is that term funding is allocated according to the highest net return subject to the bounds placed by the market limits. This version of the model is not a convex problem because, without penalties, the objective function is a linear combination of funding costs and asset returns. Therefore, the model achieves a corner solution. Periods with no issuance occur when no liabilities come due. We use this exercise as a benchmark because it displays the funding solution that would prevail without the risk adjustment derived from the preferences of the government. Figure 9 shows the U.S.-dollar term issuance decisions by funding source for every period when the penalties are set to zero. All funding is allocated to 5-year CCS up to the limit of $2 billion per period, and the periods without issuance are periods in which there was no term funding necessary, either because nothing came due or because the liquidity tier was not met during that

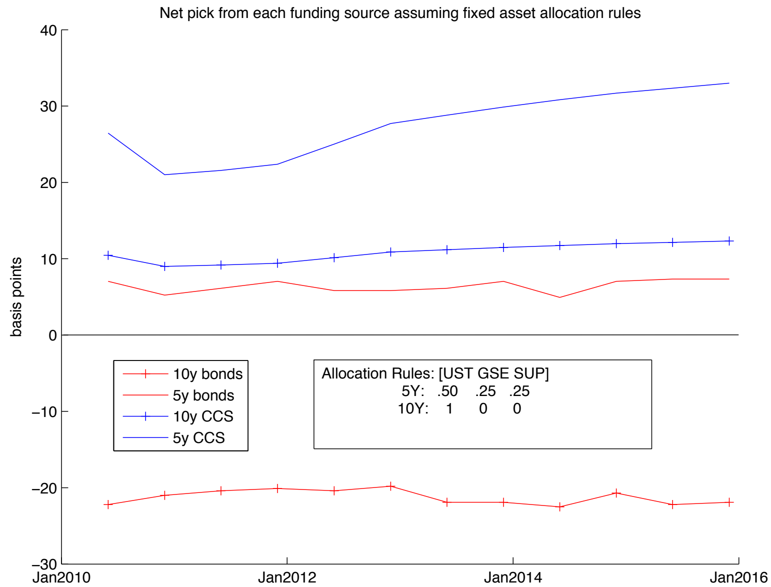

Figure 7: Forecast of the net return of CCS and global bonds for the 5- and 10-year maturities. In a net return basis (without considering risks from default and rollover), the 5-year CCS is the best funding source. The 10-year CCS and 5-year global bond follow closely at around an 8 to 10 basis point return. The 10-year global follows with a net cost at around 20 basis points.

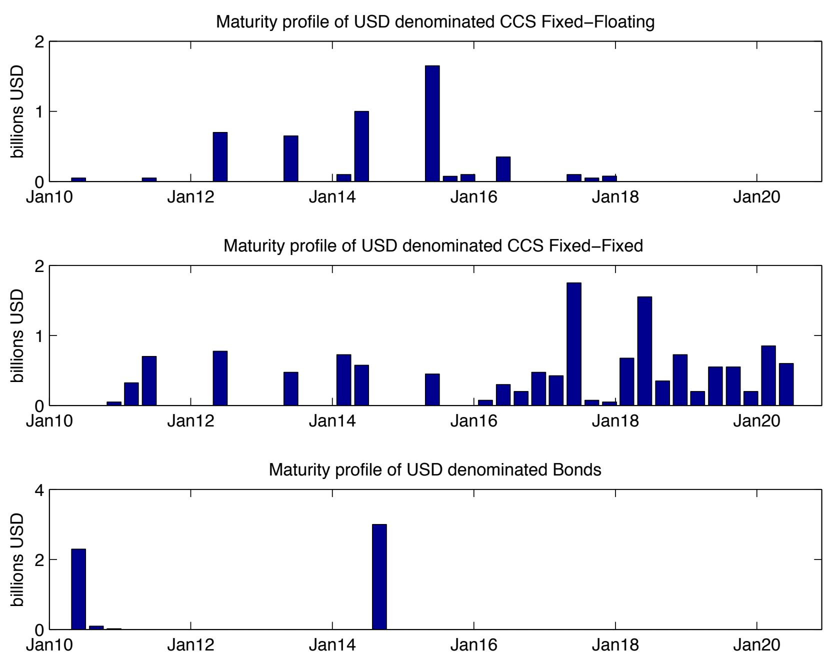

Figure 8: Actual maturity profile of the EFA at the end of 2009. The first panel shows the fix-to-floating CCS which, due to their structure, fund the liquidity tier. The second panel shows the fix-to-fix CCS. The third panel shows the global bonds and Canada bills.

period. Also, EFA size increases are set to zero, which maintains the U.S.-dollar portfolio at \$25 billion. The parameters of the penalty functions are set to zero, thus dropping out of the objective function. Figure 10 shows the resulting maturity profile, which is more skewed than the initial profile because most of the issuance is through 5-year CCS. The only remaining long-term liabilities are due to the binding market limit of \$2 billion for 5-year CCS in periods when term funding exceeded \$2 billion.

The second exercise results are shown in Figures 11 and 12. In this exercise, term funding is allocated to 5-year CCS up to the limit of \$2 billion per period. The rest is issued in 5-year bonds because the charges offset the net return advantage of the 10-year CCS, but are not large enough to offset the net return advantage of the 5-year CCS. Note, however, that with the parameterization of the credit charge function, this version of the model is a convex problem. The parameters of the penalty functions are $\theta^{5y} = 1.3$ and $\theta^{10y} = 1.6$. With this parameterization for a \$2 billion issuance of 5-year CCS, the penalty is equal to 15.5 basis points. Thus, even with the penalty, the 5-year CCS is more attractive in terms of net return than the 5-year bond, because the difference in their return is, on average, close to 20 basis points. The resulting maturity profile is completely concentrated between the 0 and 5-year tenors.

Next we take the calibrated penalties for CCS and for the skewed maturity profile discussed above. The results are shown in Figures 13 and 14. This exercise adds the maturity profile charge to the credit charges. We apply this charge only to the short-term U.S.-dollar liabilities because the effect of this penalty is through the relative changes in returns between short and long maturities. The maturity profile penalty is 20 basis points for funding quantities around \$1 billion, with a parameter $\eta^{5y} = 1.8$. Recall from the first-order condition of the model that the penalties add linearly and are subtracted from the net return of each instrument. In short, the maturity profile penalty lowers the relative value of short-term funding. Another way to do the same would be to change the portfolio weights; however, the effects through the penalty are more pronounced because the penalty function is exponential, while the portfolio weights are linear. Term funding is allocated more evenly to 5-year CCS, 10-year CCS and some 5-year global bonds. Still, 10-year bonds are not issued at all. The maturity profile is somewhat concentrated in the short term compared to the initial profile.

Another output of the model is to compute the expected asset-liability gap. Since the expected gap is computed period-by-period, this measure can be used to evaluate whether a particular funding plan would imply at any point a large deviation (e.g., ± 1 per cent of the EFA value) to the government's budgetary position. The expected gap is computed using the regression coefficients and the total amount of CCS implied by the funding decisions (see section 3.5). The independent variables are the funding cost forecasts and the total amount of CCS implied by the funding decisions. Figure 15 shows the dynamics of the asset-liability gap for the three exercises above. Note that, given the amount of CCS suggested by the exercises, the gap remains around 0.5 per cent for almost two years before returning slowly to close to zero. In each exercise, the implied share of CCS to total funding varies significantly. Yet, as argued above, the main component of the movements of the AL gap is the forecast of the Can\$/US\$ basis, not the total amount outstanding of CCS. In the benchmark exercise, the share of CCS to total funding rises from 67 per cent at the beginning to 99 per cent toward the end of the forecast period. In the exercise with the credit and maturity profile penalty, CCS remain at 68 per cent of funding.

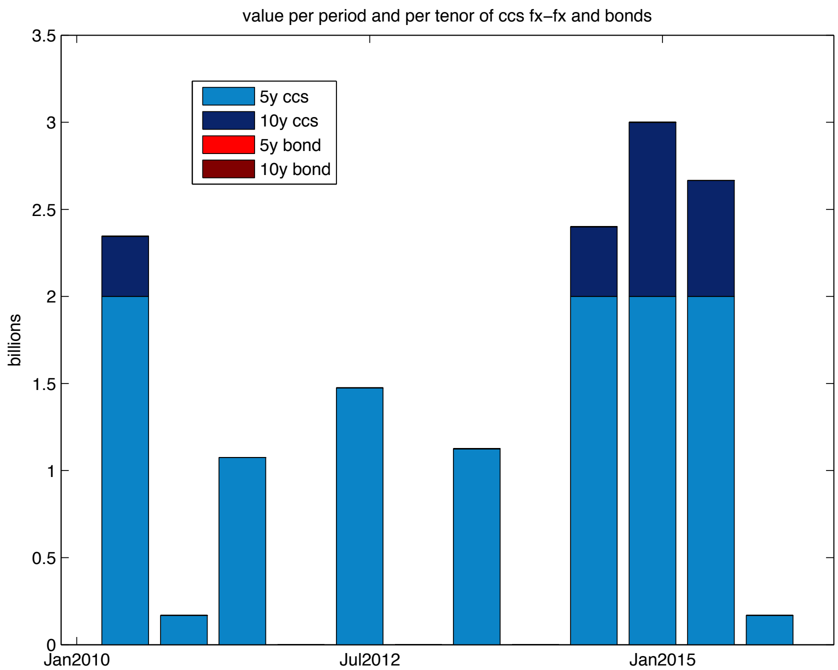

Figure 9: U.S.-dollar term issuance decisions over time for the benchmark exercise. Note that all funding is allocated to 5-year CCS up to the limit of $2 billion per period. The periods without issuance are periods in which there was no term funding necessary, either because nothing came due or because the liquidity tier was not met during that period.

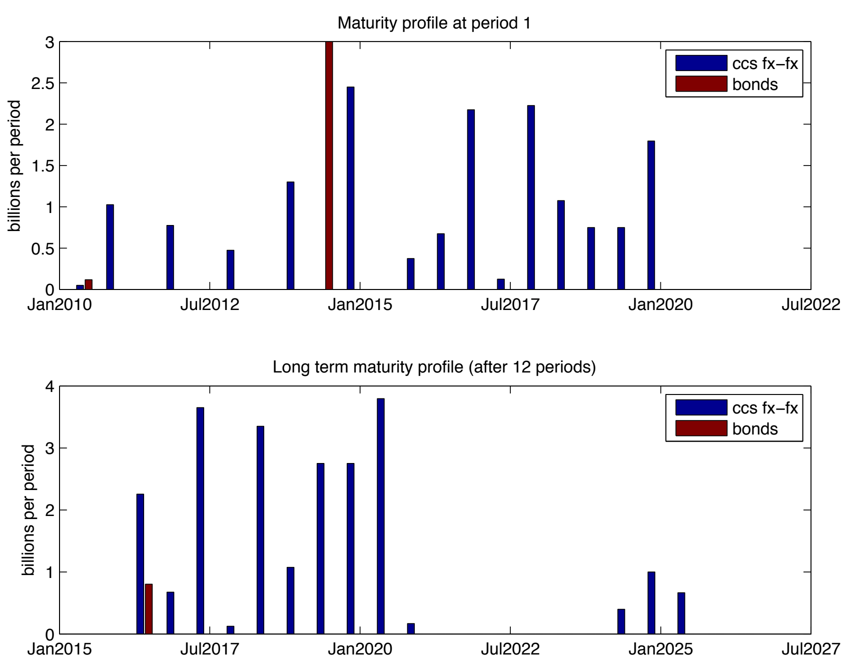

Figure 10: Resulting U.S.-dollar maturity profile of the EFA for the benchmark exercise without replacement cost or maturity profile penalties. The first panel shows the maturity profile at period 1, and the second after 12 periods. Clearly, the maturity profile becomes more concentrated in the short term.

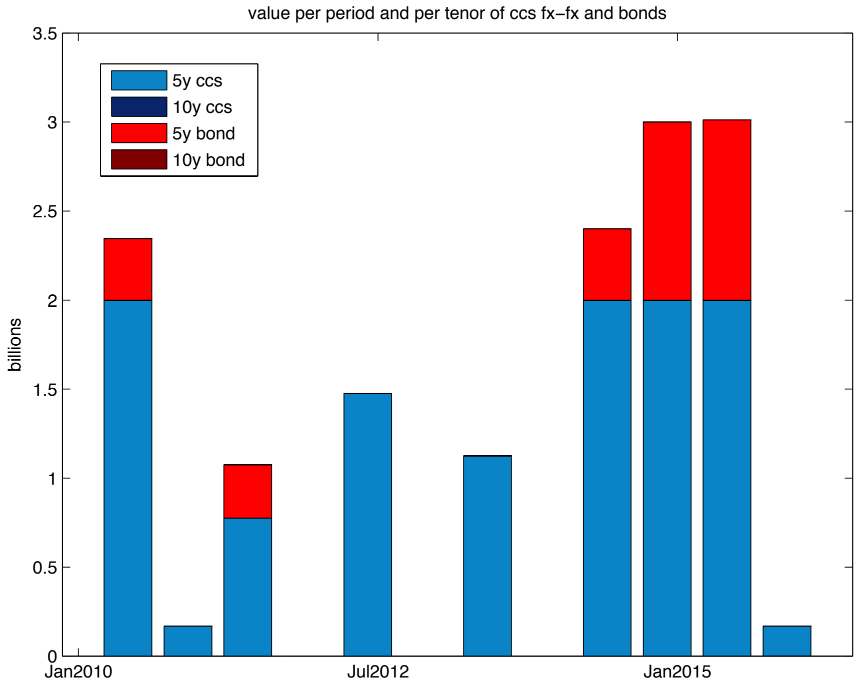

Figure 11: U.S.-dollar term issuance decisions over time for the exercise with only credit charges. Note that all funding is allocated to 5-year CCS up to the limit of $2 billion per period, except for the third period. The rest is issued in 5-year bonds, because the charges offset the net return advantage of the 10-year CCS but are not large enough to offset the net return advantage of the 5-year CCS.

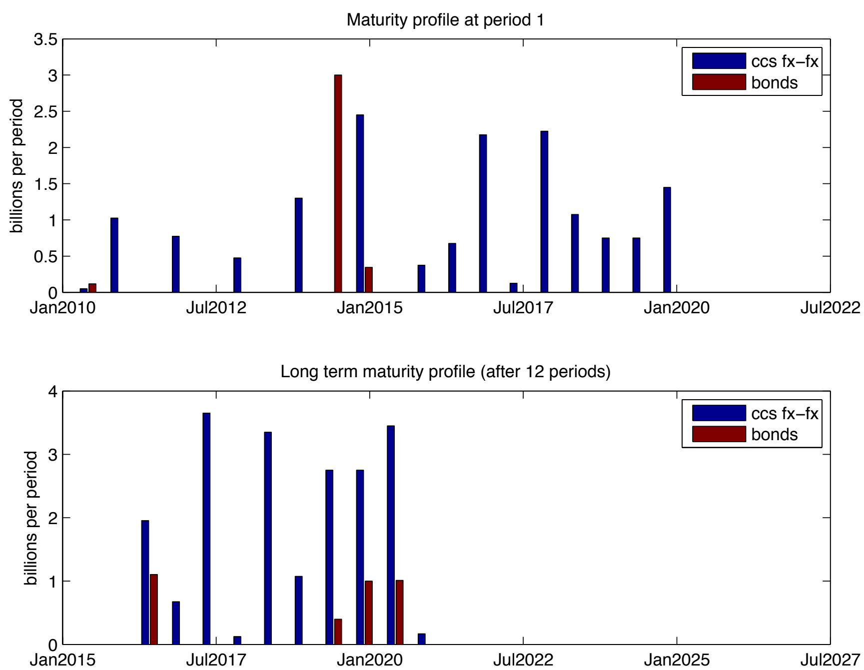

Figure 12: Resulting U.S.-dollar maturity profile of the EFA for the exercise with only credit charges. The second panel shows the maturity profile after 12 periods. The maturity profile is completely concentrated in the short term because no 10-year instrument is issued.

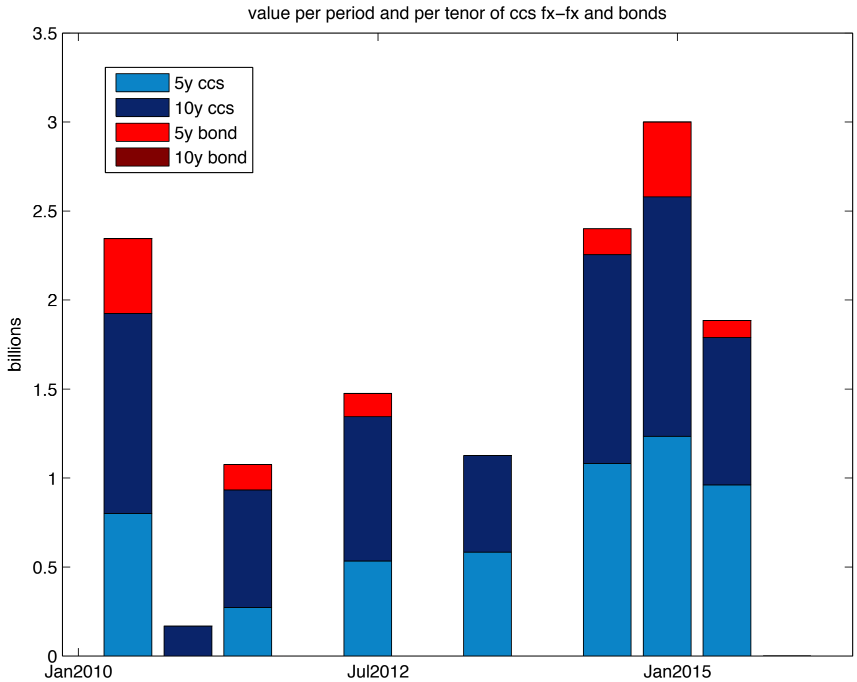

Figure 13: U.S.-dollar term issuance decisions over time for the benchmark exercise. Funding is allocated more evenly to 5-year CCS, 10-year CCS and some 5-year global bonds. Still, 10-year bonds are not issued at all.

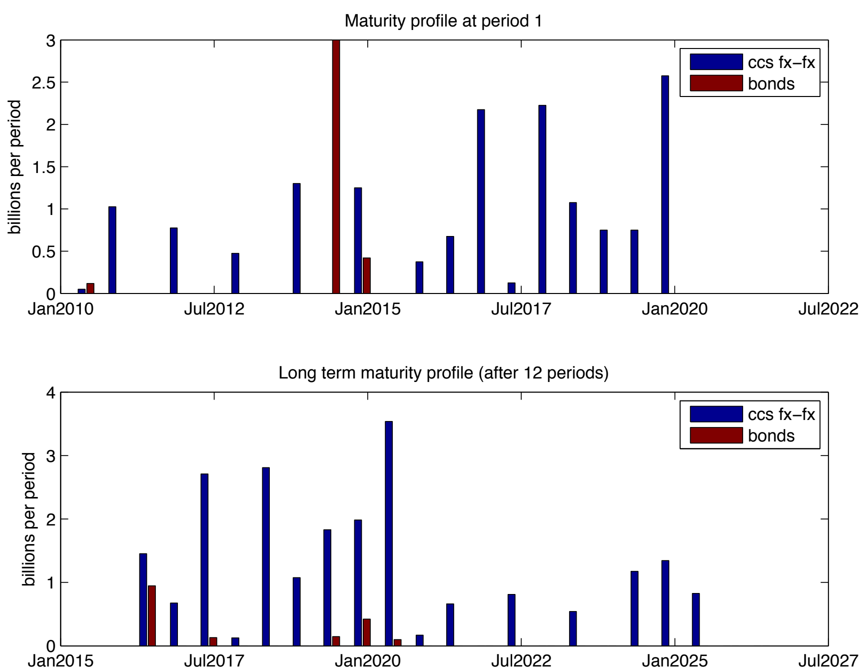

Figure 14: Resulting U.S.-dollar maturity profile of the EFA for the exercise with penalties. The second panel shows the maturity profile after 12 periods. The maturity profile is still somewhat concentrated in the short term compared to the initial profile.

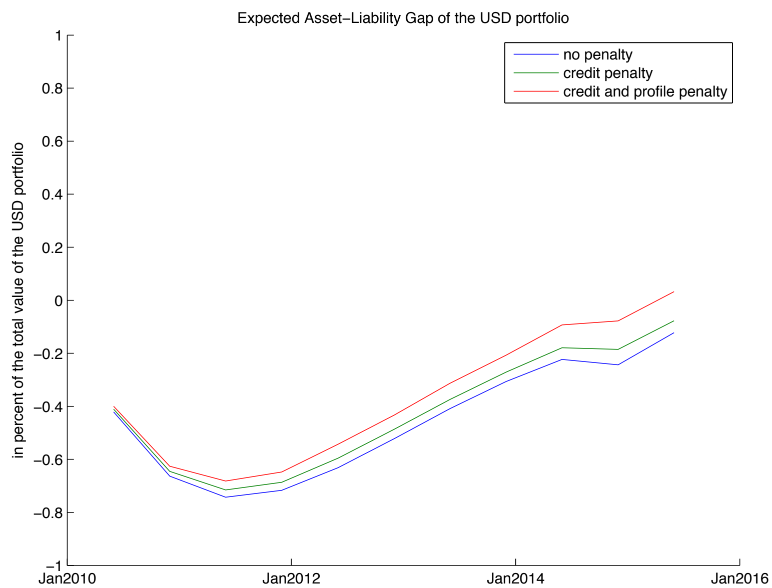

Figure 15: Implied U.S.-dollar asset-liability gap as a percentage of the U.S.-dollar EFA portfolio calculated using the regression described in section 3.5. Each line corresponds to one of the exercises shown above. The independent variables are the funding cost forecasts and the total amount of CCS implied by the funding decisions in each of the exercises. Note that, given the amount of CCS suggested by this exercise, the gap remains around 0.5 per cent for almost two years before returning slowly to close to zero for the three exercises, although they differ significantly in the implied share of CCS to total funding. In the benchmark exercise, at the end of the exercise, 99 per cent of funding is in CCS, while, in the credit and profile penalty exercise, CCS remain at 68 per cent of funding. The main component of the movements is the forecasts of the Can$/US$ basis.

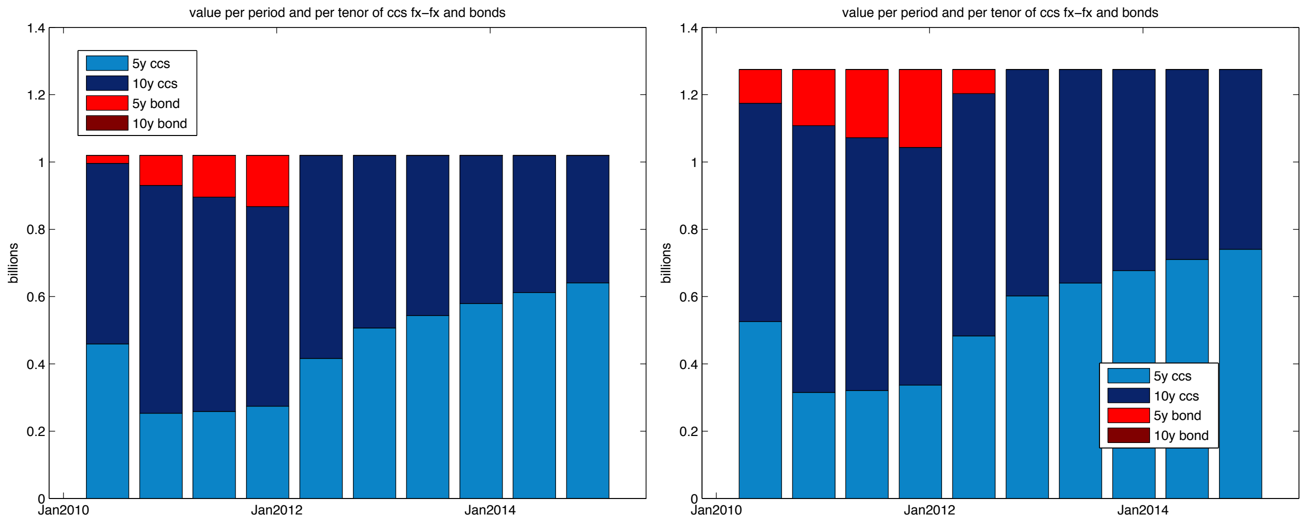

Figure 16: EFA size exercise. This plot shows the effects of the EFA size on diversification through the curvature of the penalty functions. The exercise is to compare two funding strategies of different size, all else being equal. In the left panel the issuance is \$1 billion and in the right panel the issuance is \$1.25 billion every period, while maintaining asset returns and funding costs the same as in the benchmark exercise. As the size of funding increases, the 5-year bond is more attractive because both 5- and 10-year CCS become relatively more expensive. The percentage increase in CCS is smaller than the increase in funding. The EFA maturity profile is structured to comply with liquidity tier requirements.

4.3 EFA size and diversification

The next exercise illustrates the effects of the size of the EFA on diversification. The asset returns and funding costs are the same as in the benchmark exercise. The penalty parameters are kept unchanged as well. In this exercise, the results show the curvature introduced by the exponential function of counterparty and rollover risk. The exercise is designed to force the model to issue US\$1 billion or US\$1.25 billion every period, complying with liquidity tier requirements. At the US\$1 billion level, and given the forecasts and calibrated penalties, a 5-year bond is issued in small quantities. As the size increases, the 5-year bond is more attractive, because both CCS become relatively more expensive. The share of the 5-year bond increases, while the share of the 5-year CCS falls, as well as the share of the 10-year CCS (see Figure 16). The pronounced U-shaped changes between 5- and 10-year CCS are a consequence of the narrowing of the difference of the net return between the 5-year CCS and the other funding sources (see Figure 7). The main point of this exercise is to show how the curvature of the penalty function changes the *proportion* of funding in each source, depending on the size of the EFA.

4.4 Intuition of the results and interpretation

There are three main drivers of the model. The first is the curvature introduced by the replacement cost function of CCS and the penalty from a skewed maturity profile. We assume that replacement costs are a function of outstanding amounts and are valued in terms of the expected dollar cost of replacement in case of default. The curvature changes the allocation by inducing diversification, because, in the margin, additional dollar borrowing becomes relatively more expensive (riskier) in any source and tenor.

The second driver is the asset returns and funding costs. As specified in this paper, the funding costs of bonds and the returns of the three asset classes are linear combinations of exogenous spreads and the U.S. Treasury yields. From the first three principal-component factors in the U.S. yield curve (Litterman and Scheinkman 1991), only the slope is important for the model, since the level affects both sides of the balance sheet equally. The curvature factor does not play a role because the model considers only two tenors. The slope matters because it determines the trade-off between short- and long-term sources of funding through the interaction with the penalty for skewness of the maturity profile, directly penalizing 5-year *net* returns.

The third driver is the size of the EFA. The larger the size of the EFA, on average, the larger the placements required period-by-period. This is specially true during periods of EFA increases. The larger the amount of issuance in a particular period, the larger the penalty imposed.

Judgment has to be exercised when interpreting the results, because of two features of the model. First, the period-by-period solution restricts the financing strategies available to the model. However, given the limited amount of instruments available to the manager, this may not be a tightly binding restriction. It is unlikely that the expected net return from more complicated dynamic strategies would be much larger, given that the choices for issuance are limited to 5- and 10-year instruments. On the other hand, relaxing the budget constraint assumption to allow more or less funding in a given period could significantly change the results, because, currently, the model is required to issue all of the amount coming due in a given period, while there are other periods with no issuance at all. Given the curvature of the penalty functions, the Ramsey principle implies that the expected net return would increase if the issuance could be smoother over time.

Second, the assumption regarding the market limits is that no more than $2 billion can be issued in a given instrument in a given period. However, there is also a lower limit for a given issuance of global bonds, normally at $1 billion per transaction. For the medium-term notes and European medium-term notes programs, this lower limit does not apply because small placements are common. The size of a single global bond placement is designed to take into account syndication costs. However, the model does not restrict the minimum size because putting a lower bound would *require* that amount to be placed, while the correct restriction should be that *if* there is a global bond placement, it should be larger than or equal to the limit. We avoid this type of restriction because it would introduce non-linearities into the model. Therefore, when the model suggests a small placement of global bonds, it can be interpreted as an upper bound of suggested issuance. Also, it can be considered intertemporally, meaning that small placements over the course of several periods could be lumped together in a single placement that satisfies the transaction size requirements in the market.

4.5 Euro exercises

For completeness, we perform a set of exercises in the euro part of the EFA. All assets and liabilities in this exercise are quoted in euros and use the actual maturity profile of the EFA at the end of 2009. We construct euro forecasts of asset returns and funding costs in a way that is equivalent to the U.S.-dollar forecasts. However, our assumptions for the asset-allocation rules are different from the U.S.-dollar case,

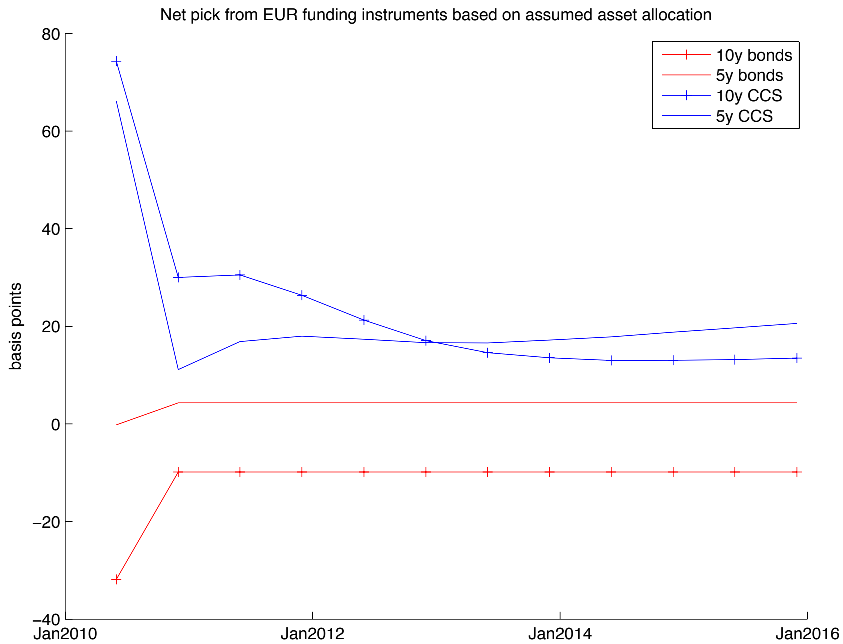

Figure 17: Forecast of the net return of euro CCS and euro global bonds for the 5- and 10-year maturities. On a net return basis (without considering risks from default and rollover), the 5- and 10-year CCS are the best funding sources. The asset-allocation rule used to compute net return in this case is for the 5-year tenor [0.15, 0.27, 0.58] and the 10-year tenor [0.26, 0.45, 0.29].

because of the larger availability of euro instruments. The portfolio weights by tenor are the following: $\begin{pmatrix} \alpha^{5y} & \beta^{5y} & \gamma^{5y} \end{pmatrix} = \begin{pmatrix} 0.15 & 0.27 & 0.58 \end{pmatrix}$ and $\begin{pmatrix} \alpha^{10y} & \beta^{10y} & \gamma^{10y} \end{pmatrix} = \begin{pmatrix} 0.26 & 0.45 & 0.29 \end{pmatrix}$. Figure 17 shows the expected net returns from each funding source based on these parameter values. Another difference with the U.S.-dollar exercise is that the euro part of the EFA has no liquidity tier requirements. Regarding penalties, we maintain the same penalty parameters from the last U.S.-dollar exercise.

Figure 18 shows the issuance profile and the resulting maturity profile. The results show that most funding is allocated to 10-year CCS, except later in the sample, when the net return of CCS falls bellow the net return of 5-year CCS. At around 1 billion euros, the penalty for the profile renders both the 5-year bond and 5-year CCS less attractive than the 10-year instruments.

5 Concluding Remarks

This paper presents a simple return maximization model to guide the liability issuance decisions of the EFA. The model computes the trade-off between the cost of the debt and two sources of risk: rollover and counterparty risks. The main drivers of the model are the penalties on the counterparty risk from cross-currency swap contracts and the maturity profile risk. The solution is period-by-period but, in general, a dynamic model may not provide much greater benefits given the institutional constraints on active management of liabilities of the EFA. A simple extension to the model would be to endogenize the choice of portfolio weights, α, β

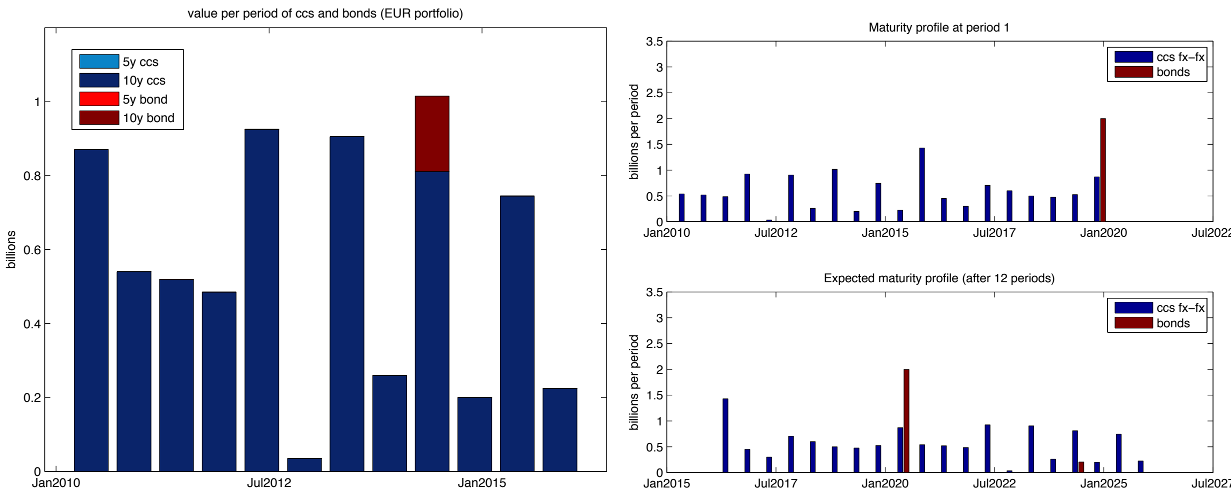

Figure 18: Term issuance decisions over time for the euro portfolio exercise and resulting maturity profile. Note that all funding is allocated to 10-year CCS, except in one period in which a 10-year global bond is issued. The first panel shows the maturity profile at period 1 and the second after 12 periods.

and γ, that determine the net return of financing strategies. This can be done numerically by allowing the optimizer routine to choose these parameters to maximize the same objective function. In this way, the model would consider the optimal asset allocation.

References

Bolder, D. 2006. "Modelling Term-Structure Dynamics for Risk Management: A Practitioner's Perspective." Bank of Canada Working Paper No. 2006-48.

Campbell, J. Y. 2006. "Household Finance." *Journal of Finance* 61 (4): 1553–1604.

Claessens, S., J. Kreuser and R. Wets. 2000. "Strategic Risk Management for Developing Countries: An Application to Colombia." Presented at the second Inter-regional Debt Management Conference in Geneva, Switzerland, April 2000.

Cochrane, J. H. 1999. "Portfolio Advice for a Multifactor World." *Federal Reserve Bank of Chicago Economic Perspectives* 23 (3): 59–78.

Department of Finance Canada. 2009. "Report on the Management of Canada's Official International Reserves, April 1, 2008 - March 31, 2009." Government of Canada. Available at http://www.fin.gc.ca/toc/2009/oir09-eng.asp.

Detemple J. and M. Rindisbacher. 2008. "Dynamic Asset Liability Management with Tolerance for Limited Shortfalls." *Insurance Mathematics and Economics* 43 (3): 281–94.

Duffie, D. and M. Huang. 1996. "Swap Rates and Credit Quality." *Journal of Finance* 51 (3): 921–49.

Duffie, D. and K. J. Singleton. 2003. *Credit Risk: Pricing, Measurement, and Management.* Princeton: University Press.

Geyer, A. and W. T. Ziemba. 2008. "The Innovest Austrian Pension Fund Financial Planning Model InnoALM", *Operations Research* 56 (4): 797–810.

Litterman, R. and J. Scheinkman. 1991. "Common Factors Affecting Bond Returns." *Journal of Fixed Income* 1 (1): 54–61.

Markowitz, H. 1952. "Portfolio Selection." *Journal of Finance* 7 (1): pp. 77–91.

Missale, A. 1999. *Public Debt Management.* Oxford: University Press.

Mulvey, J. M. and W. T. Ziemba. 1998. "Asset and Liability Management Systems for Long-Term Investors: Discussion of the Issues." In *Worldwide Asset and Liability Modeling*, 3–38. Edited by W. T. Ziemba and J. M. Mulvey. Cambridge, UK: Cambridge University Press.

Romanyuk, Y. 2010. "Asset-Liability Management: An Overview." Bank of Canada Discussion Paper No. 2010–10.

Wachter, J. 2010. "Asset Allocation." National Bureau of Economic Research Working Paper No. 16255.

Zenios, S.A. and W. T. Ziemba. 2006. *Handbook of Asset and Liability Management*, Amsterdam: North-Holland.